

during initial configuration of the network, thereby providing each node within the network with sufficient information to compute a complete map of the entire network. Thereafter, global LSUs are only transmitted upon the expiration of a period of time specified by a periodic timer. Between global LSU transmissions, non-global LSUs are transmitted. Typically, the TTL value of each non-global LSU is set to a value smaller than that of the size of the network so that they do not propagate throughout the entire network. Upon expiration of the period of time specified by the periodic timer, each switching node again transmits a global LSU.

Because global LSUs are only transmitted within this hazy-sighted routing scheme on a periodic basis, during certain periods of time various nodes within a network implementing this routing protocol may lack up-to-date information regarding the exact location of every other node in the network. Thus, although nodes may have received sufficient information to compute an up-to-date map of their surrounding region (as determined by the TTL value of the most recent non-global LSU), their understanding of the location of or best path to distant nodes (i.e., nodes outside of their horizon line) may be based on out-of-date information (as determined by the most recent global LSU).

Hazy-sighted routing thus allows information about distant nodes to be inexact, such that a switching node always knows how to get a packet closer to a destination node, but may not always know the details of the best path to this destination node. Once a transmitted packet has been forwarded to a node that is closer to the destination node, more information about this path is provided, and so on at the next closest node until the packet eventually arrives at the destination node. Inasmuch as the number of topological changes that might occur within the time specified by the periodic timer is likely to be greater than one, this periodic timing limitation serves to reduce the number of LSUs generated, thereby limiting the amount of traffic overhead promulgated within the network. Hazy-sighted routing thus sacrifices routing accuracy in favor of reduced link-state overhead.

As with the traditional link-state approach to routing described above, the so-called “hazy-sighted” routing approach suffers from similar scalability, performance and reliability concerns. For example, it is still necessary for the routing table to contain information about each and every node. Global LSUs are essential to providing such information. Thus, as discussed above, the use of global LSUs limits scalability, network performance and reliability.

In addition, although the amount of routing accuracy sacrificed by hazy-sighted routing schemes to achieve the desired reduction in link-state traffic overhead is an acceptable trade-off for networks composed of low-capacity links (such as ad-hoc networks formed exclusively of wireless connections), this tradeoff is problematic in networks composed at least in part of high-capacity links. For example, in a network composed of numerous high-capacity links, the amount of bandwidth that is saved by reducing the TTL of an LSU is relatively minor in comparison with the total bandwidth of each high-capacity link in the network. Thus, when hazy-sighted routing schemes are adapted for use in connection with networks composed at least in part of high-capacity links, the relatively minor traffic overhead efficiencies that are achieved are generally outweighed by the resulting losses in routing accuracy.

Accordingly, there exists a need for a system and method capable of enabling nodes within an ad-hoc network to seamlessly communicate with adjacent nodes, distant nodes and a wider network (such as the Internet) so long as physical connectivity is maintained with at least one other node. There

also exists a need for a system and method capable of scaling beyond the size limitations of traditional ad-hoc networks, while minimizing any potential decreases in network performance and reliability. In addition, there exists a need for a system and method capable of adjusting the TTL of a packet based on the capacity of the links over which this packet will travel. Preferably, such a system and method would provide significant improvements in scalability, application performance and overall network connectivity.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an exemplary physical topology of a network capable of implementing a routing protocol.

FIG. 2 illustrates an exemplary physical topology of a network capable of implementing a no-sight routing protocol.

FIG. 3 illustrates an exemplary physical topology of a network in which the travel of data may be limited based on variable translucency time-to-live values.

FIG. 4 depicts an exemplary process flow for performing no-sight routing.

FIG. 5 depicts an exemplary process flow for limiting the travel of data via variable translucency time-to-live values.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

I. System Overview

a. No Sight Routing

FIG. 1 is a block diagram illustrating an exemplary physical infrastructure of a system **100** for implementing a no-sight routing protocol. FIG. 2 illustrates an exemplary physical topology of a sub-network **125** implementing an exemplary no-sight routing protocol with predetermined propagation limits.

Exemplary system **100** generally comprises, among other things, nodes **102, 104, 106, 108, 110, 112, 114, 116, 118** and **120**. Nodes **102-120** are connected to one another via connections **122**, which may include any number of connections recognized in the art, including, for example, wires, wireless communication links, fiber optic cables, etc. Nodes **102-120** connected together via connections **122** collectively form sub-network **125**.

In general, nodes **102-120** represent connection terminals within exemplary sub-network **125**. In some embodiments, a protocol operating on a network above that of sub-network **125** distinguishes between nodes **102-120** based on their packet-forwarding capabilities. For example, in some embodiments a protocol operating on a network above that of sub-network **125** recognizes oval-shaped nodes **102, 104, 106, 112, 114, 118** and **120** as representing “hosts” (i.e., nodes which only forward originating packets, as will be known to those of skill in the art) and rectangular-shaped nodes **108, 110, and 116** as representing “routers” (i.e., nodes which forward/route non-originating packets). This host/router distinction is not, however, made within sub-network **125**. Within sub-network **125**, all nodes are viewed as being directly connected; i.e., any node can send data to any other node.

According to certain embodiments, one or more of nodes **102-120** collectively forming exemplary sub-network **125** may be a mobile node. Generally speaking, a mobile node is a device whose location and point of attachment to exemplary sub-network **125** may frequently change. Examples of mobile nodes include cellular telephones, handheld devices, PDAs, and portable computers.